January 1978

SPECIFICATION OF INTERNETWORK TRANSMISSION CONTROL PROGRAM

TCP
Version 3

Vinton G. Cerf
Advanced Research Projects Agency

Jonathan B. Postel
Information Sciences Institute

January 1978

Information Sciences Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, California 90291

(213) 822-1511

INFORMATION SCIENCES INSTITUTE

UNIVERSITY OF SOUTHERN CALIFORNIA

-- Messages from file: [PARC-MAXC]<SHOCH>MESSAGE.IX;1
-- WEDNESDAY, FEBRUARY 15, 1978 08:19:50-PST --

Date: 14 FEB 1978 1602-PST
From: CERF at USC-ISI
Subject: Proposals for new Internet and TCP header formats
To: [ISIE]<Postel>INTERNET-Working-Group.List:,
To: [ISIE]<Postel>TCP-Working-Group.List:

Mail from USC-ISIE rcvd at 14-FEB-78 1511-PST
Date: 14 Feb 1978 1514-PST
From: CERF
Subject: test
To: cerf

< CERF, INTERNETHEADER.NLS.11, >, 14-Feb-78 15:09 VGC
;;;;.H1Sw=on;.H2Sw=on;.H3Sw=on;

Became IEN #26

Internet Notebook Section 2.3.2.1
IEN No. YY

Vint Cerf
ARPA
14 Feb 1978

2.3.2.1 A Proposed New Internet Header Format

As a result of the recent TCP and internet working group meetings at the end of January, the notion of an internet layer of protocol has become somewhat clearer. The basic internet service is datagram oriented, but also accommodates the fragmentation of datagrams at gateways, reassembly taking place at the destination gateway associated with the destination host. Of course, fragmentation and reassembly of datagrams within a network or by private agreement between the gateways of a network is also allowed since this is transparent to the internet protocols and the higher-level protocols. This transparent type of fragmentation/reassembly is termed "network-dependent" fragmentation and is not dealt with further here.

At the meetings, the point was also raised that addressing my be confined to the internet header at least for the current host/host protocols such as TCP and the two undefined but planned protocols: datagram protocol (DGP) and real-time protocol (RTP). This strategy will better support multi-protocol synchronization when this is required for multi-media teleconferencing.

In the format proposal which follows, an attempt has been made to reduce the impact on header length of the new features of the internet protocol. It is proposed that some flexibility be given up in favor of reducing header sizes or processing requirements. Figure 1 illustrates the proposed new internet header format.

The proposed new internet format only allows two versions ("VER" field) of the internet header to co-exist. This seems sufficient to deal with any reasonable transition period during which one header is being phased out. Once the internet system is in operation on a regular basis, changes to the header format should be rare since such changes would affect all host and gateway software.

Inter-Office Memorandum

To Distribution

Date July 10, 1978

From L. Stewart

Location Palo Alto

Subject June 1978 TCP meeting notes and MIT AI Lab visit.

Organization SSL

XEROX

TCP 6/78

MIT/LCS

Filed on: [Maxc1] < LStewart > TCP678.memo

John Shoch and I attended a meeting of the TCP working group on June 15-16, 1978, at MIT LCS. The meeting was attended by people from MIT, BBN, SRI, DTI (Digital Technology Inc.), LLL, ARPA, DOD, ISI, and Xerox. Vint Cerf was chairman aided by Jon Postel. The TCP (Transmission Control Program), implements an internetwork Host-Host protocol based on a proposal by Cerf and Kahn (A Protocol for Packet Network Intercommunications. IEEE Transactions on Communications, COM-22, May 1974).

Status of the TCP World

The current stage in the development of TCP is a move towards bringing up TCP servers and users around the Arpanet, but TCP remains clouded by technical details and continuing arguments about its proper functions.

Implementations of TCP Version 2.5 now exist for the LSI-11 under MOS, for the PDP-11 under Unix, for the IBM-360 under TSO, and both Bcpl and assembler implementations for Tenex. Implementations for RSX-11 and Multics are expected soon. Meanwhile, TCP Version 4 has been more or less finalized (by the June meeting) and some conversion efforts have started. Quite reasonably enough, under the circumstances, there is a 'version' field in the TCP header.

TCP/Telnet is in regular use between LSI-11 *Terminal Interface Units* on the Packet Radio network and the Arpanet hosts which support TCP servers. MIT LCS apparently now plans to use TCP rather than its home grown 'DSP' as the end-to-end protocol on the not-yet-operational LCSNet; and, according to the DOD representative, a sizeable fraction of potential military data net users have decided to adopt TCP. This group is a source of a great deal of pressure to "Keep the header size down". A TCP based FTP is currently being designed.

TCP Discussions

Practically all of the changes made to the TCP specification at the meeting concern the interface between TCP and an underlying Internet Datagram layer. Over the past six months, many of the early features of TCP which gave it such a 'Monolithic Protocol' appearance have been moved out of the TCP and into the Internet layer. The most important of these are the *Type of Service*, *Internet Addressing*, and *Fragmentation* mechanisms. TCP now specifes a reliable byte stream given a raw packet transport mechanism. *In brief, TCP is now much closer to the spirit of Pup than before.* It seems as though a good deal of the reason for the monolithic nature of the early TCP was that various groups were unable to agree on a clean layering of protocol.

While most of the byte-stream and connection handling aspects of TCP have stabilized, there was still considerable confusion over the meaning and utility of the *Urgent* mechanism. Eventually, the group agreed that Urgent is an out-of-band signal that means 'Read faster, something important is up ahead.' Urgent is something like the Pup BSP Interrupt. *There is no TCP mechanism corresponding to Pup BSP Marks.* The other major TCP internal discussion was about the precise interpretation of *letter* boundaries and the ways in which they interact with Urgent; this issue is not fully resolved.

On the second day, a great deal of time was spent arguing about the internet address format. Perhaps the discussion would have been better placed at a meeting of the Internet working group, but the two groups do have considerable overlap. After many hours, the group compromised on the usual [net, host, socket] format with fixed field lengths. The socket field was left in what might best be described as the gray zone between the datagram layer and the TCP. It became apparent at the meeting that most of the TCP community thinks of demultiplexing taking place first by protocol and then by socket. *This is the reverse of PUP.* It should be noted that the fixed field addressing represents a backing away from a rather free format address recommended only six months ago. Under the former plan an address consisted of a sequence of (address segment size, address segment) pairs.

MIT AI Lab notes

While in Boston, I paid a visit to the MIT AI Lab (ninth floor) and was given a tour by Dave Weinreb.

Chaosnet

The Chaos network is a single coax very much Ethernet-like system under development at the AI Lab. The cable operates at 8 megabits, the transceivers operate through cable taps (currently tee connectors), and the interfaces provide full packet buffering. The system is operational and connects the MIT-MC ITS system, the Lisp machine, and possibly a few others. A fairly lengthly (26 page) in-progress paper by Dave Moon describing the philosophy and protocols of the net may be found on [Maxc]<LStewart>Chaord.

The fully buffered nature of the chaosnet interface is due to the high cable speed and the very long interrupt latency times expected. Currently, the only interface designed runs on a PDP-11 style Unibus and the maximum packet length is around 1000 bits. The MIT-MC interface hangs on the KL-10 front end PDP-11 while the Lisp machine has a Unibus of its own.

LISP Machine

The MIT AI lab has constructed a microprogrammed machine to run LISP at high speeds. The processor is Schottky TTL and consists of about 9 square feet of wire-wrap board. The I/O system and main memory are additional. The basic machine includes a four plane 512 square bitmap display, but no disk. The group plans to utilize file servers on the Chaosnet. Excepting microcode, the entire system is written in Lisp. One machine is operational and a second under construction. Initially the lab plans to build about 40 machines for about $40,000 each.

c: Parc: David. Boggs, Ed Taft, John Shoch, Larry Stewart, Burt Sutherland, Ted Strollo
SDD: Ron Crane, Yogen Dalal, Bob Metcalfe, Hal Murray, Jim White
ASD: Doug Brotz, Steve Butterfield
Rochester: Jack Wheeler

After 2/78 meetings in CA

Draft Internetwork Protocol Specification

Version 2

1. INTRODUCTION

The Internet Protocol is designed for use in interconnected systems of computer communication packet-switched networks. The internet protocol provides for transmitting segments of data from sources to destinations, where sources and destinations are identified by variable length addresses. The internet protocol also provides for fragmentation and reassembly of long segments, if necessary, for transmission through "small packet" networks.

1.1. History

This protocol has been developed as one result of the ARPA sponsored internetwork experiments program. The history until January 1978 is the history of the host-to-host protocol TCP.

The first publication of the ideas on which TCP is based was a paper in the IEEE Transactions on Communications by Cerf and Kahn in 1974 [1]. Later that year a protocol specification was published by a group led by Cerf at Stanford University [2]. A second specification was prepared in 1976 by a group led by Postel at SRI for the Defense Communication Agency for the AUTODIN II network [3]. In 1977 Cerf, at ARPA, prepared a substantial revision of the TCP specification [4]. Recently Postel revised Cerf's revision to distinguish the internet aspects from the host-to-host aspects [5].

1.2. Scope

The internet protocol is specifically limited in scope to provide the functions necessary to deliver a package of bits (an internet segment) from a source to a destination over an inconnected system of networks. There are no mechanism to promote relibility, flow control, sequencing, or other services commonly found in host-to-host protocols.

The protocol is intended to be utilized in gateways that interconnect sets of networks.